REMARKS

Reconsideration of the present application is respectfully requested. No claims have been amended. No claims have been canceled or added. No new matter has been added.

Claim Rejections

Independent claims 1, 8, 13, 19, 24, 30, 32 and 34 stand rejected under 35 U.S.C. § 102(a) based on Tehloh (US 2003/0014523). Applicants respectfully traverse the rejections.

Claim 1 recites:

> 1. A method for mirroring data on a first storage server and a second storage server, the method comprising:
> **queuing write commands at the first storage server between consistency points, the write commands being to write data corresponding to a file system of the first storage server to a local mass storage device coupled to the first storage server;**
> **at a start of each consistency point, sending the write commands to the local mass storage device and to a remote mass storage device coupled to the second storage server;**
> updating memory blocks of the local and remote mass storage devices based on the write commands; and
> at an end of each consistency point constructing a representation to reference each memory block of the local mass storage device that is in use to represent the file system; and sending at least a portion of the representation to the second storage server. (**Emphasis added**).

In contrast, Tehloh does not teach or suggest the above emphasized limitations. Regarding the limitation of queuing write commands at the first storage server between consistency points, the write commands being to write data corresponding to a file system of the first storage server to a local mass storage device coupled to the first storage server, the Examiner contends that Tehloh's Figure 1 discloses data replicating in which data and commands are queued in the local replication facility (office action mailed on 11/22/2006, page 3). Applicants studied Figure 1 and relevant sections regarding Figure 1, but found no discussion regarding queuing write commands

between consistency points. Rather, blocks 36 and 40 in Tehloh's Figure 2 and paragraphs 46 and 48 disclose that data to be written, instead of the corresponding write commands, is queued and transmitted to the remote mass storage device (e.g., the mirror). As disclosed in paragraph 46, data package forwarded to the remote site includes information that identifies a storage location, such as a volume path. The data package, however, is completely different from a set of write commands. Although paragraph 43 of Tehloh discloses logging write commands to a local storage device, paragraphs 43 and 50 of Tehloh further explain that the logging occurs in the event of a disruption in the remote mirroring process, such as an outage, not between consistency points such as recited in claim 1. As discussed in paragraph 6 of the present application, in a system which handles large volumes of client requests, it may be impractical to save data modifications to the mass storage devices every time a write request is received from a client. The reason for this is that mass storage device accesses tend to take a relatively long time compared to other operations. Therefore, the source file server may instead hold write requests in memory temporarily and save the modified data to the mass storage devices periodically, such as every 10 seconds or at whatever time interval is appropriate. The event of saving the modified data to the mass storage devices is called a "consistency point". In contrast, Tehloh contains no discussion regarding the concept of consistency point.

Regarding the limitation of at a start of each consistency point, sending the write commands to the local mass storage device and to a remote mass storage device coupled to the second storage server, the Examiner does not even specifically point out which section of Tehloh is considered to read on this limitation. As discussed above, Tehloh's Figure 2 and paragraphs 46 and 48 disclose that data to be written, instead of the corresponding write commands, is transmitted to the remote mass storage device (e.g., the mirror).

Thus, at least for the foregoing reasons, Tehloh does not teach or suggest the above emphasized limitations of claim 1. Therefore, claim 1 and all claims which depend on it are patentable over Tehloh.

Independent claims 8, 19, 24 and 30 each recite limitations similar to those discussed above for claim 1. For similar reasons, claims 8, 19, 24, 30 and the claims which depend on them are also patentable over Tehloh.

Claim 13 recites the limitation of receiving block-level write commands to update memory blocks of a local mass storage device coupled to the second storage server, based on changes to the file system (of a first storage server). This limitation of claim 13 is recited from the perspective of the second storage server, but is similar to the limitation of sending the write commands to a remote mass storage device coupled to the second storage server, recited in claim 1. Thus, for similar reasons, claim 13 and all claims which depend on it are patentable over Tehloh. Claim 32 recites a similar limitation to that discussed above for claim 13. Thus, for similar reasons, claim 32 and all claims which depend on it are patentable over Tehloh.

Claim 34 recites:

> 34. A method, comprising:
> in a first storage server, constructing a representation to reference each memory block used to store a file system, the memory blocks being part of a first set of mass storage devices comprising at least one mass storage device coupled locally to the first storage server; and
> **sending at least a portion of the representation to a second storage server, said portion comprising information to allow reconstruction of the entire representation by the second storage server so that the second storage server has a representation of memory blocks of a second set of mass storage devices comprising at least one mass storage device coupled locally to the second storage server used to store the file system.**
> (Emphasis added)

The Examiner contends that Tehloh's paragraph 46 teaches or suggests the above emphasized limitation (Office Action mailed on 11/22/06, page 5). The Examiner alleges that

paragraph 46 discloses the local data replication facility sending information in the data package including volume path and the data to the remote data replication facility so that the data can be stored in the second set of mass storage devices. Applicants respectfully disagree.

As disclosed in paragraph 46, the data package sent to the remote site includes data to be replicated and information identifying a storage location. Neither does paragraph 46 disclose, nor does the nature of the data package suggest, that such data comprises information to allow reconstruction, on a remote storage server, of a representation referencing each memory block used to store a file system of a first storage server. Thus, at least for the foregoing reasons, Tehloh does not teach or suggest the above emphasized limitations for claim 34. Claim 34 and all claims which depend on it are patentable over Tehloh.

If any additional fee is required, please charge Deposit Account No. 02-2666.

Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: February 22, 2007

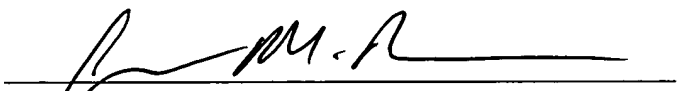

Jordan M. Becker
Reg. No. 39,602

Customer No. 48102
12400 Wilshire Boulevard
Seventh Floor
Los Angeles, CA 90025-1030
(408) 720-8300